

5 **A SYSTEM AND PROCESS FOR DISCOVERY OF NETWORK-CONNECTED
DEVICES**

BACKGROUND

10 Technical Field:

The invention is related to discovery among network-connected devices, and more particularly to a system and process for a discovering electronic device to discover the presence and network address of a discoverable electronic 15 device connected to the same network and located in the same delimited space, without employing the network.

Background Art:

20 Communication between electronic devices connected via a local network is a common requirement in a variety of applications, such as found in ubiquitous computing environments and some teleconferencing systems. For example, ubiquitous computing revolves around extending computational activities beyond the current desktop model and into the environment. In future homes and 25 offices, access to computing should be as natural as access to lighting. Users should not be required to go to a special place (i.e., the desktop) to interact with the computer. Rather, the computer should be available to interface with the user anywhere in the home or office (or more generally anywhere in an arbitrarily large environment), through whatever set of devices is available, be they fixed or 30 carried by the user. In order to affect this ubiquitous computing environment the various endpoint components and computing devices (hereinafter referred to collectively as devices) need to communicate with each other. One typical way

this is accomplished is via communication over a conventional local computer network. This can be done either using wired or wireless connection schemes.

In the context of teleconferencing, meetings are an important part of everyday life for many workgroups. A variety of live teleconferencing systems are available commercially to enhance and archive these meetings. Both recorded meetings and live video conferences require audio-visual capturing equipment. One such system generally referred to as a Distributed Meeting (DM) system provides high quality multi-party conferencing and recording of meetings, as well as rich browsing of archived meetings enhanced through a number of analysis techniques. DM systems typically have multiple cameras that simultaneously capture different sub-events occurring in a space where an event occurs, such as a meeting room, to give a rich experience to local and remote meeting participants. These capture devices are all synchronized to provide an integrated, scalable system and method for two-way conferencing, broadcasting, recording and viewing meetings or other events. In particular, a DM system's cameras can include a 360-degree camera centrally positioned to monitor in 360 degrees the space in which the event occurs; a remote view camera positioned so as to capture a view of event participants in the meeting room; a presenter view camera positioned so as to capture a view of the front of the meeting room where a presenter is typically presenting; and a whiteboard capture camera positioned so as to capture strokes written on a whiteboard. A DM system also typically has one or more microphone devices or microphone arrays that simultaneously capture audio data that is synchronized with the image data recorded by the various cameras. A given microphone device or array can be used to determine the direction from which an event participant is speaking by using sound source localization. The 360 degree camera, presenter view camera and remote view camera all can have microphone arrays associated with them. A DM system can also include a projector for projecting meeting or other materials onto a screen and a graphics capture device is used to capture graphics data used in the meeting, such as electronic slides that are part of an

electronic presentation. A monitor or television may also be included to display one or more remote participants or other meeting/event materials. Further, a DM system also can have a kiosk that is used to control broadcast, recording and playback of the meeting or other event. In addition, at least one server or other 5 suitably configured computer that is capable of storing and distributing the captured data to one or more clients in network connection with the server can be included.

For live meetings using a DM system, the system broadcasts the 10 multimedia meeting streams to participants, who use a network for low-latency duplex voice and data communication. The meetings can also be recorded and viewed on-demand. These viewing and archiving features can be useful not only to remote users who are not at the site of the meeting, but also to participants 15 who do attend the meeting. To facilitate the viewing and archiving of a live meeting by local participants, the participant will typically employ a portable computing device, such as a notebook/laptop computer or even a handheld computing device. For example, a local meeting participant may want to use their notebook computer to record portions of the meeting, such as the audio and/or what is written on a whiteboard. To accomplish this task, the network 20 addresses for the endpoint component associated with the microphone array used to capture the audio portion of the meeting, and the whiteboard camera (or other conventional whiteboard capture device), must be known to or discovered by the notebook computer.

Thus, in a DM scenario there are many different devices that must 25 communicate with each other. One effective way of accomplishing this task is to connect all the devices to a conventional local computer network. Here again, this can be done either using wired or wireless connection schemes.

However, in the foregoing scenarios the various devices may be new to the environment (e.g., a participant's notebook computer) or a fixed device may be activated and brought on-line to the network after having been deactivated for a period of time (e.g., the various cameras employed in a DM system). As a result, the various components and computing devices will not know of each others existence. Accordingly, a key to establishing communications between devices over a network is for the devices to discover each other. Once a connection is established, a device can control or be controlled by other devices, send or receive event signals, and transfer data between applications.

10

Universal Plug and Play (UPnP) is a standard framework that enables discovery between components and devices connected to a network, and is the typical method for establishing communications. In general, this framework typically requires a device to multicast its address over the network. Other components or devices seeking communication with the multicasting device would then use the address to establish communications. However, this method of discovery presents serious security issues. For example, it allows any device connected to the network to find and communicate with any other device on the network regardless of their relative locations. This can be a problem in the context of a distributed meeting because it is often desired to restrict control of the conference room endpoint components to participants actually in the conference room. For example, it may be desired by the attending participants to control what portions of a meeting are recorded via the endpoint components. If persons outside the conference room can readily access and remotely control these endpoint components then the desired control cannot be exercised. Further, unless potentially elaborate security measures are taken, unauthorized access to a device is a possibility using the standard UPnP discovery methods. As a result of these concerns, among others, businesses having local networks often disable the multicast feature. Thus, there is no way for devices to discover each other and so the advantages of UPnP cannot be realized.

15

20

25

30

Thus, there is a need for alternate discovery techniques that establish communications between endpoint components and computing devices over a network in a manner that localizes the communications links to a prescribed space, such as a conference room. In addition, there is a need for alternate discovery techniques that can be combined with UPnP to allow UPnP to work even in environments with the multicast capability disabled.

SUMMARY

The present invention is directed toward a system and process for discovery between endpoint components and computing devices (hereinafter referred to collectively as devices) residing in the same delimited space, such as conference room, in cases where the devices are connected to a common local network. The present invention allows a device in the space to discover the presence and network address of another device in the space, which then allows the first device to negotiate with the discovered device to facilitate the transfer of data and other communications over the common network. However, unlike prior schemes this discovery is accomplished without the use of the common network, and so is not affected by the disablement of the network multicast feature typically employed for discovery purposes. In addition, discovery is limited to discovering devices residing in the space also containing the discoverable devices.

More particularly, in one embodiment of the present discovery system and process, the network address assigned to the discoverable device located within the delimited space is transmitted by that device. This signal is not transmitted via the common network. In addition, it is transmitted in a manner that substantially limits its reception to the delimited space, thereby precluding the aforementioned localization issues associated with discovery via a network. The signal can also be periodic in nature, if desired, so that it is received by new devices coming into the space. The discovering device receives the signal and

uses the received network address to establish communications via the common network with the discoverable device that transmitted the address. The aforementioned signal can further include a signature expressly indicating the type of electronic device that sent the network address. It is noted that the fact 5 that such a signal is sent by a device also inherently indicates it is present in the space and is accessible via said common network. Whenever the signal includes the discoverable device's signature, the discovering device receiving the signal would only initiate communications with that device if the discoverable device is of the type that it is desired contact. In one version of this first 10 discovery embodiment, the signature is in the form of an identifier that distinguishes the associated discoverable electronic from all other discoverable electronic devices in the space.

In another embodiment of the present discovery system and process, a 15 discovering device transmits a signal, which requests the network address assigned to a discoverable device in the space to be transmitted to the discovering device. This signal is also not transmitted over the network, is of a type that substantially limits its reception to the delimited space, and can be periodic in nature or not as desired. A discoverable device receiving the request 20 signal transmits a reply signal having the same characteristics and which includes data representing the requested network address. Additionally, the request signal can include data specifying the type of discoverable device it is desired to obtain the network address for, whereby only a device of the specified type transmits a reply signal. Alternately, if the device type is not included in the 25 request, the reply signal transmitted by the discoverable device can include the aforementioned device signature, thereby allowing the discovering device to determine whether the replying device is of a type it wants to establish communications with over the network.

30 In yet another embodiment of the present discovery system and process, a discoverable device within the space transmits a signal including the

aforementioned device signature. This transmission can be periodic so as to be received by discovering devices newly entering, or coming online within, the delimited space. The discovering device receives the signature signal, and if the discoverable device is of the type it is desired to establish communications with,
5 transmits the aforementioned request signal. As in the previous embodiment, the discoverable device receives the request and transmits a reply signal with its network address. The discovering device then receives the reply signal and establishes communications with the discoverable device over the common network. Here again, the signature, request and reply signals are not transmitted
10 via the network and are transmitted in a manner that substantially limits their reception to the delimited space.

The foregoing embodiments of the discovery system and process can be implemented using infrared (IR) or audio signals. In the case of IR-based discovery, each discoverable and discovering device includes an IR transmitter and/or receiver as needed to transmit or receive the aforementioned signals, depending on which discovery embodiment is employed. It is noted that the IR transmitters are powerful enough to extend throughout the delimited space.
15

20 In the case of audio-based discovery, the discoverable and discovering devices include a microphone and loudspeaker, as appropriate, for receiving and transmitting audio signals respectively. The audio signal transmitted can be audible or inaudible. Inaudible audio signals are especially appropriate if the transmitted signal is periodic, so as not to distract people in the delimited space.
25 If an inaudible signal is employed, it can be a low amplitude audio signal inaudible to humans that has data embedded in it in the form a conventional audio watermark. In this case the receiving device obtains the desired data by ascertaining the content of the audio watermark. If an audible signal is employed, the data is encoded in the signal using any appropriate conventional
30 audio encoding scheme (e.g., DTMF encoding), and the receiving device obtains the desired data by decoding the signal using the appropriate decoding scheme.

It is further noted that regardless of if IR or audio signals are employed, the signal can be encrypted for security purposes using any appropriate conventional encryption scheme. In such a case the device receiving the signal would decrypt it prior to obtaining the desired data. Another security measure that can be implemented to ensure only discovering devices present in the delimited space can obtain the network address of a discoverable device also located in the space, involves including a confirmation actuator on the discoverable device. This actuator must be physically activated by person within the delimited space before the discoverable device will transmit a signal containing its network address. In embodiments where a discoverable device's signature is transmitted in a separate signal from the network address, it can be required that the aforementioned actuator be activated prior to the signature signal being transmitted as well.

15

In addition to the just described benefits, other advantages of the present invention will become apparent from the detailed description which follows hereinafter when taken in conjunction with the drawing figures which accompany it.

20

DESCRIPTION OF THE DRAWINGS

The specific features, aspects, and advantages of the present invention will become better understood with regard to the following description, appended 25 claims, and accompanying drawings where:

FIG. 1 is a diagram depicting a general purpose computing device constituting an exemplary system for implementing the present invention.

FIG. 2 is a flow chart diagramming an exemplary process for implementing one embodiment of the present discovery technique using IR signals.

5 FIG. 3 is a flow chart diagramming an exemplary process for implementing another embodiment of the present discovery technique using audio signals.

10 FIG. 4 is a flow chart diagramming an exemplary process for implementing yet another embodiment of the present discovery technique using either IR or audio signals.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

15 In the following description of preferred embodiments of the present invention, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. It is understood that other embodiments 20 may be utilized and structural changes may be made without departing from the scope of the present invention.

1.0 The Computing Environment

25 Before providing a description of the preferred embodiments of the present invention, a brief, general description of a suitable computing environment in which the invention may be implemented will be described. Fig. 1 illustrates an example of a suitable computing system environment 100. The computing system environment 100 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use 30 or functionality of the invention. Neither should the computing environment 100

be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment 100.

5 The invention is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

10 The invention may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

15 With reference to Fig. 1, an exemplary system for implementing the invention includes a general purpose computing device in the form of a computer 110. Components of computer 110 may include, but are not limited to, a processing unit 120, a system memory 130, and a system bus 121 that couples various system components including the system memory to the processing unit 120. The system bus 121 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus

using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus.

Computer 110 typically includes a variety of computer readable media. Computer readable media can be any available media that can be accessed by computer 110 and includes both volatile and nonvolatile media, removable and non-removable media. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computer 110. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of the any of the above should also be included within the scope of computer readable media.

The system memory 130 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 131 and random access memory (RAM) 132. A basic input/output system 133 (BIOS), containing the basic routines that help to transfer information between elements within computer 110, such as during start-up, is typically stored in ROM 131. RAM 132 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 120. By way of example, and not limitation, Fig. 1 illustrates operating system 134, application programs 135, other program modules 136, and program data 137.

10

The computer 110 may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, Fig. 1 illustrates a hard disk drive 141 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 151 that reads from or writes to a removable, nonvolatile magnetic disk 152, and an optical disk drive 155 that reads from or writes to a removable, nonvolatile optical disk 156 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 141 is typically connected to the system bus 121 through a non-removable memory interface such as interface 140, and magnetic disk drive 151 and optical disk drive 155 are typically connected to the system bus 121 by a removable memory interface, such as interface 150.

25
30
The drives and their associated computer storage media discussed above and illustrated in Fig. 1, provide storage of computer readable instructions, data structures, program modules and other data for the computer 110. In Fig. 1, for example, hard disk drive 141 is illustrated as storing operating system 144, application programs 145, other program modules 146, and program data 147.

Note that these components can either be the same as or different from operating system 134, application programs 135, other program modules 136, and program data 137. Operating system 144, application programs 145, other program modules 146, and program data 147 are given different numbers here to illustrate that, at a minimum, they are different copies. A user may enter commands and information into the computer 110 through input devices such as a keyboard 162 and pointing device 161, commonly referred to as a mouse, trackball or touch pad. Other input devices (not shown) may include a joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 120 through a user input interface 160 that is coupled to the system bus 121, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 191 or other type of display device is also connected to the system bus 121 via an interface, such as a video interface 190. In addition to the monitor, computers may also include other peripheral output devices such as speakers 197 and printer 196, which may be connected through an output peripheral interface 195. Of particular significance to the present invention, a microphone array 192, and/or a number of individual microphones (not shown) are included as input devices to the personal computer 110. The signals from the microphone array 192 (and/or individual microphones if any) are input into the computer 110 via an appropriate audio interface 194. This interface 194 is connected to the system bus 121, thereby allowing the signals to be routed to and stored in the RAM 132, or one of the other data storage devices associated with the computer 110. Further, an infrared (IR) transceiver 198 is included as an input/output device to the personal computer 110. The signals to and from the IR transceiver 198 are input into and output from the computer 110 via an appropriate IR interface 199, which is connected to the system bus 121. It is noted that rather than an IR transceiver, an IR transmitter or IR receiver could be employed instead, as desired.

The computer 110 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 180.

The remote computer 180 may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 110, although only a memory storage device 181 has been illustrated in Fig. 1. The logical connections depicted in Fig. 1 include a local area network (LAN) 171 and a wide area network (WAN) 173, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the computer 110 is connected to the LAN 171 through a network interface or adapter 170. When used in a WAN networking environment, the computer 110 typically includes a modem 172 or other means for establishing communications over the WAN 173, such as the Internet. The modem 172, which may be internal or external, may be connected to the system bus 121 via the user input interface 160, or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 110, or portions thereof, may be stored in the remote memory storage device. By way of example, and not limitation, Fig. 1 illustrates remote application programs 185 as residing on memory device 181. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

2.0 System And Process For Discovering Networked Endpoint Components And Computing Devices

The exemplary operating environment having now been discussed, the remaining part of this description section will be devoted to a description of the program modules embodying the invention. Generally, the system and process

according to the present invention involves discovery between endpoint components and computing devices (hereinafter referred to collectively as devices) in a delimited space, such as conference room, in cases where the devices are connected to the same local network. As described previously, 5 discovery is the first step in order for two devices to communicate. More particularly, the present invention allows a device in a space to discover the presence and network address of another device in the space, which then allows the first device to negotiate with the discovered device for data transfer, control, and so on.

10

The foregoing discovery system and process can be implemented using infrared (IR) and audio. Each of these scenarios will be described in the sections to follow.

15

2.1 Infrared-Based Discovery

20

In general, infrared (IR) based discovery in accordance with the present system and process requires that any device that is to be discovered in a space include an IR transmitter. These IR devices are conventional and used in a variety of computing devices. As such no detailed description of them need be provided herein. In addition, any device that is to discover another device in the space must have an IR receiver. In cases where a device can act as either a discovering or discoverable device, it would require both an IR transmitter and receiver.

25

30

Referring to Fig. 2, an example of implementing the present IR-based discovery technique for a first discovery embodiment is provided. This involves first having the device that is to be discovered transmits an IR signal containing its signature and network address via its IR transmitter using conventional methods (process action 200). This signal can be periodic in nature. Meanwhile, another endpoint component or computing device in the space

listens for the signature and address with its IR receiver (process action 202). When it is determined that this information has been received (process action 204), the discovering device starts negotiations for data transmission (e.g., UPnP) via the common local network (process action 206). It is noted that the
5 signature of a device can indicate several things such as the type of device, that it is a device connected to the same local network as the discovering device, and simply that the device is present in the space. While the latter two items can be assumed to be inherent in the fact that a device is transmitting a signature within the space, the first item can require that some unique identifier be used that
10 distinguishes the device from all other discoverable devices in the space.
Ideally, the unique identifier representing a device's signature would be made as small as possible, and a prescribed list of identifiers each identifying its corresponding device would be known to all the devices capable of discovering other devices in the space.

15 Referring now to Fig. 3, an example of the foregoing IR-based discovery in the aforementioned Distributed Meeting (DM) setting might involve two endpoint components, i.e., a 360 degree camera 300 and a whiteboard camera 302) within a meeting space establishing communications with each other to synchronize data capture and the like. In addition, a portable PC device 304 (e.g., notebook computer, tablet-type PC, and so on) is placed in the space by a meeting participant the purpose of discovering the 360 degree camera 300 in order to obtain and record portions of the meeting such as the audio and video streams captured by the camera 300. First, assume the aforementioned 360
20 degree camera 300 needs to establish communications over the local network 314 with the whiteboard camera 302 to synchronize data capture. For this purpose, the white board camera 302 would include an IR transmitter 306, and a network connection 316 for connecting to the network 314 via a network hub
25 322. The 360 degree camera 300 needs to know the IP address of the whiteboard camera 302 in order to establish communications with it over the local network 314. As indicated previously, this is accomplish by having the
30

whiteboard camera 302 periodically transmit its signature and network address in an signal output by its IR transmitter 306. The 360 degree camera 300, which is equipped with an IR receiver 308, and a network connection 318 for connecting to the network 314 via the network hub 322, listens for a signal from the

5 whiteboard camera 302 (as identified by its transmitted signature) and obtains the network address of the device. Communications between the cameras 300, 302 is then established by the 360-degree camera over the local network 314 in a conventional manner via their network connections 316, 318 and the hub 322.

As the 360 degree camera 300 is also a discoverable device, it includes an IR transmitter 310 as well. The PC device 304 needs to discover the 360 degree

10 camera 300 in order to establish communications with it over the network to capture the desired meeting data. Thus, the PC device 304, which is equipped with at least an IR receiver 312, and a network connection 320 for connecting to the network 314 via a network hub 322, listens for a signal from the 360 degree

15 camera 300 (as identified by its transmitted signature) and obtains the network address of the device. Communications between the camera 300 and the PC device 304 are then established by the computer over the local network 314 in a conventional manner using their network connections 318, 320 and the hub 322.

Thus, IR-based discovery is accomplished in a seamless manner with no user

20 intervention required.

It is noted that the aforementioned IR transmitter should in some circumstances exhibit more power than a typical IrDA transmitter, which is limited to a 1 m range. More particularly, the IR transmitter's power should be strong enough to extend throughout the space. For example, the IrDA transceiver (Part No. GP2W2001YK) produced by Sharp Microelectronics of the Americas located in Camas, Washington, which has a 5 meters transmit range, could be employed in a conference room setting as most conference rooms are not larger than 5 meters in any direction.

IR-based discovery has important advantages over standard techniques, such as aforementioned network-based multicast scheme, which solve the problems addressed previously. For instance, IR-based discovery is inherently local. IR signals typically stay within a room, as they do not penetrate walls, go through glass, and are not easily reflected off of most wall surfaces. Thus, in a conference room setting, the signals could not be received outside the room. In addition, the use of the present IR-based discovery technique does not require any infrastructure changes. For example, as mentioned earlier the multicast feature in many local networks is disabled, which prevents discovery with protocols like UPnP. However, the present IR-based discovery technique allows devices within a delimited space to discover each other's network address, without the use of multicasting over the network. Thus, communication between devices can be established without any change to the network settings. In this way, IR-based discovery can be combined with UPnP to allow UPnP to work even in environments with the multicast feature disabled. Even if the multicast feature were not disabled on the local network, the use of the IR-based technique would be advantageous as it solves the aforementioned locality issue by only allowing devices present in the delimited space to discover each other.

20 2.2 Audio-Based Discovery

In general, audio-based discovery in accordance with the present system and process requires the devices to be discovered, as well as the devices performing the discovery, in a delimited space, to each include a loudspeaker and microphone (which can be one of the aforementioned microphone arrays employed in DM systems). It is noted that most portable PC devices already have these items, and they can be readily incorporated into any endpoint component.

30 Referring to Fig. 4, an example of implementing the present audio-based discovery technique using a second discovery embodiment is presented. This

involves first having the device that is to be discovered transmit a low amplitude audio signal that is inaudible to humans, and which has an embedded signature, using its loudspeaker (process action 400). This signal can be periodic in nature. Another endpoint component or computing device in the space samples audio
5 with its microphone looking for the signature (process action 402). When it is determined that the signature has been received (process action 404), the discovering device transmits an audible audio signal including a request for the network address of the device being discovered using its loudspeaker (process action 406). The device being discovered then receives the request via its microphone and transmits an audible audio signal in reply that includes its network address in encoded form (process action 408). The discovering device next starts negotiations for data transmission (e.g., UPnP) via the common local network to which the devices are connected using the received address (process
10 action 410). It is noted that the signature transmitted by the device being discovered would have the same attributes as the one employed in the previously-described IR-based discovery technique.
15

The audio signal transmitted by the discoverable device that includes its signature can be an audible audio signal instead, if desired. Further, the audio signal transmitted by the discovering device to request a network address can alternately be an inaudible signal similar to the periodically transmitted signature. Similarly, the reply signal transmitted by the device being discovered can be of
20 the inaudible type.

An inaudible audio signal is implemented in one embodiment of the present audio-based discovery technique by encoding the desired data (e.g., the device's signature, a request for a network address or the network address itself) via a conventional watermark technique. In general, an audio watermark involves altering the frequency spectrum of an audio signal to embed a pattern.
25 This pattern represents the data being encoded. In the context of producing an inaudible watermark encoded audio signal in accordance with the present
30

system and process, the audio signal is low-amplitude noise, which is inaudible to a human, and which has been altered as described above to embed the desired data. These alterations do not change the inaudible nature of the signal.

The use of inaudible signals has an advantage of not being a distraction to

5 people in the space. This is particularly true in the case of the periodic signal.

As for audible audio signals, these can be implemented using any conventional audio-based encoding method, such as via the use of DTMF encoding. An advantage of using an audible signal is that it can have a relatively short duration in comparison to an inaudible signal employing watermarking techniques, even though providing the same amount of embedded data. For example, a typical network address can be encoded in an audible signal with a duration of 320 ms or less depending on the encoding used, whereas it would require an inaudible watermarked signal several seconds to encode the address.

10 In addition, since the audio signals used to request an address and provide one in reply would occur infrequently, they would not present a significant distraction to people in the area even if they were audible. To the contrary, the use of an audible signal to request and provide a network address would alert the user to the fact that communication is being established between two devices in a space. Thus, in one preferred version of the present audio-based discovery technique, an inaudible periodic audio signal is used to transmit the signature of a discoverable device employing watermarked encoding, while the audio signals used to request a network address and provide the address in reply employ encoded audible signals.

15

20 Referring now to Fig. 5, an example of the foregoing audio-based discovery in the aforementioned Distributed Meeting (DM) setting might involve a meeting participant using their portable PC device 500 (e.g., notebook computer, tablet-type PC, and so on) to record portions of the meeting such as the audio and video streams captured by the aforementioned 360 degree camera 502. To accomplish this task, the participant's portable PC device 500 must first discover

the camera's local network address. The camera 502 and portable PC device 500 are each equipped with a loudspeaker 504a, 504b, microphone 506a, 506b, and network connection 508a, 508b for connecting to the network 510 via a network hub 512. The camera 502 would periodically transmit a low amplitude 5 inaudible audio signal that contains its signature encoded as an audio watermark. In one version of the present audio-based discovery technique, the signature was embedded as a 1-bit watermark in an inaudible audio signal, which is transmitted once every second. As indicated previously, the audio signal is made inaudible to humans so that it does not present a distraction in the 10 meeting space. The participant's portable PC device 500 listens for a signal from the 360 degree camera 500 by recording about 1 second of audio and decoding any audio watermarks contained in the recorded block. If the 360 degree camera's signature (which is known to the PC device 500) is found in the recorded block of audio, the PC device transmits the aforementioned audible 15 request for the camera's network address. As the request is of short duration (e.g., less than 500 ms), the fact that it is audible to humans will not present a significant distraction. The 360 degree camera 502 receives the audible request and replies with an audible signal encoded with its network address using a prescribed encoding scheme (e.g., DTMF encoding). Here again the reply is of 20 short duration (e.g., less than 500 ms) and so will not constitute a distraction to the meeting participants even though it is audible. The PC 500 then decodes the audible signal received from the camera 502, thereby obtaining the network address of the device. Communications between the PC device 500 and the camera 502 are then established by the PC device over the local network 510 in 25 a conventional manner using their network connections 508a, 508b and the hub 512. Thus, as with IR-based discovery, the present audio-based discovery technique is accomplished in a seamless manner with no user intervention required.

30 Like IR-based discovery, audio-based discovery has important advantages over standard techniques, such as aforementioned network-based

multicast scheme, which solve the problems addressed previously. For instance, audio-based discovery is inherently local. Audio amplitudes decay rapidly as a function of distance from the source. The audio signals employed in the present audio-based discovery technique typically will not travel beyond the confines of the defined space, as they do not significantly penetrate walls, and if high frequency signals are used, they will be attenuated when reflected off of most wall surfaces. Thus, in a conference room setting, the signals would not be significantly detectable outside the room. In addition, as in the case of IR-based discovery, the use of the present audio-based discovery technique does not require any infrastructure changes. For example, as mentioned earlier the multicast feature in many local networks is disabled, which prevents discovery with protocols like UPnP. However, the present audio-based discovery technique allows devices within a delimited space to discover each other's network address, without the use of multicasting over the network. Thus, communication between devices can be established without any change to the network settings. In this way, audio-based discovery too can be combined with UPnP to allow UPnP to work even in environments with the multicast feature disabled. Additionally, even if the multicast feature were not disabled on the local network, the use of the audio-based technique would be advantageous as it solves the aforementioned locality issue by only allowing devices present in a delimited space to discover each other.

It is noted that the discovery scheme described previously in connection with the description of the present IR-based technique could also be implemented using audio signals having the characteristics described in this section. Similarly, audio signals as described above in connection with the present audio-based technique as implemented in the second discovery embodiment could be replaced with IR signals containing the same data.

30 2.3 Additional Features

Some additional features of the present discovery system and process applicable to both IR-based and audio-based techniques will now be described.

2.3.1 Discovering Device Initiated Discovery

5

While the foregoing discovery techniques involved a discoverable device initiating discovery by transmitting its signature (and in some cases its network address as well) within a delimited space, an alternate discovery embodiment would be for the discovering device to transmit a request for a network address instead. Referring to Fig. 6, in this alternate scenario, the aforementioned request signal would be transmitted first (process action 600). This could be made on a periodic basis or not, and could identify what particular device it is desired to obtain an address for in the request signal. The signals can be either IR or audio. In the case of the audio-based discovery technique, if the request is made periodic in nature, the use of an inaudible watermark encoded signal would be less distracting to people in the space. However, an audible signal could be used if such a distraction is not an overriding concern or the request is not periodic.

20

The discoverable device would reply to the request with a network address as before (process action 602). The discovering device then starts negotiations for data transmission (e.g., UPnP) via the common local network to which the devices are connected using the received address (process action 604).

25

If the request is general, such that replies from all discoverable devices in the space are solicited, then the reply signal would include both the device's network address and its signature. In this way the requesting device can ascertain what type of device is replying, thereby allowing the requesting device to select which device or devices it wants to establish communications with over the local network. If the request was specific as to the type of device, then the

reply signal would not need to contain a device signature. However, as before, the mere fact that the request is answered by a device is indicative that the device is present in the space and that it is available for communications over the local network.

5

2.3.2 Encrypted Signals

As an added security measure one or more of the signals transmitted by a discoverable device can be encrypted using conventional encryption techniques, such as DES encryption. In this way, the presence of a discoverable device in a space and/or its network address could only be ascertained by a discovering device possessing the ability to decrypt the signals. This would allow access to certain devices to be limited to persons previously granted access rights.

10

2.3.3 Confirmation of Physical Presence

Yet another added security measure that can be implemented in the present discovery system and process for situations where a user wants to establish communications with another device in the space via their portable computing device involves confirming the physical presence of a user within the space before a discoverable device would respond to a request for its network address. One way of accomplishing this is to include a confirmation button, switch or the like on the discoverable device. A user wanting to establish communications with a discoverable device would proceed as described previously for either the present IR-based or audio-based discovery techniques, except that a discoverable device will only transmit its network address if the aforementioned confirmation button has been activated on the device ahead of time. Thus, a user must physically activate the confirmation button on the discoverable device in order to obtain a signal containing the device's network address. In this way the user's physical presence within the space is confirmed.

20

25

30